

Faster Longest Common Extension Queries in Strings over General Alphabets

Paweł Gawrychowski*, Tomasz Kociumaka†, Wojciech Rytter‡ and Tomasz Walen‡

Institute of Informatics, University of Warsaw
[gawry,kociumaka,rytter,walen]@mimuw.edu.pl

Abstract

Longest common extension queries (often called longest common prefix queries) constitute a fundamental building block in multiple string algorithms, for example computing runs and approximate pattern matching. We show that a sequence of q LCE queries for a string of size n over a general ordered alphabet can be realized in $\mathcal{O}(q \log \log n + n \log^* n)$ time making only $\mathcal{O}(q + n)$ symbol comparisons. Consequently, all runs in a string over a general ordered alphabet can be computed in $\mathcal{O}(n \log \log n)$ time making $\mathcal{O}(n)$ symbol comparisons. Our results improve upon a solution by Kosolobov (Information Processing Letters, 2016), who gave an algorithm with $\mathcal{O}(n \log^{2/3} n)$ running time and conjectured that $\mathcal{O}(n)$ time is possible. We make a significant progress towards resolving this conjecture. Our techniques extend to the case of general unordered alphabets, when the time increases to $\mathcal{O}(q \log n + n \log^* n)$. The main tools are difference covers and the disjoint-sets data structure.

1 Introduction

While many text algorithms are designed under the assumption of integer alphabet sortable in linear time, in some cases it is enough to assume general alphabet. A general alphabet can be either ordered, meaning that one can check if one symbol is less than another, or unordered, meaning that only equality of two symbols can be checked. Many classical linear-time string-matching algorithms (e.g. Knuth-Morris-Pratt, Boyer-Moore) work for any unordered general alphabet. Recently, a linear-time algorithm for computing the leftmost critical factorization in such model was given [11]. On the other hand, algorithms related to detecting repetitions usually need $\Omega(n \log n)$ equality tests [18], and an on-line algorithm matching this bound is known [13].

In this paper we consider the longest common extension problem (LCE, in short) in case of general ordered and unordered alphabets. The goal is to preprocess a given word w of length n for queries $\text{LCE}(i, j)$ returning the length of the longest common factor starting at position i and j in w . Such queries are often a basic building block in more complicated algorithms, for example in computing runs [1, 2] as well as in approximate string matching [15].

For integer alphabets of polynomial size, one can preprocess a given string in linear time and space to answer any LCE query in constant time. Preprocessing space can be traded for query time [4, 5] and generalizations to trees [3] and grammar-compressed strings [9, 10, 16, 19] are known. The situation is more complicated for general alphabets. If the alphabet is ordered, then of course we can reduce it to $[1..n]$ by sorting the characters in $\mathcal{O}(n \log n)$ time and preprocess the obtained string in linear time and space to answer any LCE query in constant time. However this increases the total preprocessing time to $\mathcal{O}(n \log n)$.

*Work done while the author held a post-doctoral position at Warsaw Center of Mathematics and Computer Science.
†Supported by Polish budget funds for science in 2013-2017 as a research project under the ‘Diamond Grant’ program.
‡Supported by the grant NCN2014/13/B/ST6/00770 of the Polish Science Center.

For unordered alphabet the situation is even worse, because the reduction would take $\mathcal{O}(n^2)$ time. A natural question is hence how efficiently we can answer a collection of such queries given one by one (on-line), where we measure the preprocessing time plus the total time taken by all the queries.

It is known that if we can perform on-line $\mathcal{O}(n)$ LCE queries for a given word of length n in total time $T(n)$ making $\mathcal{O}(n)$ symbol comparisons, then we can compute all runs in $\mathcal{O}(n+T(n))$ time making only $\mathcal{O}(n)$ symbol comparisons. An algorithm with $T(n) = \mathcal{O}(n \log^{2/3} n)$ time was recently presented by Kosolobov [14], who posed the existence of a linear-time algorithm as an open question. Much earlier, Breslauer [6] asked in his PhD thesis whether an easier task of square detection (equivalently, checking if a word has at least one run) is possible in linear time in the comparison model. In this paper we make a significant progress towards answering both questions by giving a faster algorithm with $T(n) = \mathcal{O}(n \log \log n)$.

Our result For a given string of length n over a general ordered alphabet, we can answer on-line a sequence of q LCE queries in $\mathcal{O}(q \log \log n + n \log^* n)$ time making $\mathcal{O}(q + n)$ symbol comparisons. In particular, a sequence of $\mathcal{O}(n)$ queries can be answered in $\mathcal{O}(n \log \log n)$ time. Consequently, all runs in a string over a general ordered alphabet can be computed in $\mathcal{O}(n \log \log n)$ time making $\mathcal{O}(n)$ symbol comparisons. For a general unordered alphabet we answer q LCE queries in $\mathcal{O}(q \log n + n \log^* n)$ time, still making $\mathcal{O}(q + n)$ symbol comparisons.

Overview of the methods At a very high level, our approach is similar to the one used by Kosolobov. We first show how to calculate $\min(\text{LCE}(i, j), t)$ efficiently, where $t = \text{polylog } n$. Then we use a difference cover to sample some positions in the text. Using “short” queries, we can efficiently construct a sparse suffix array for these sampled positions, which in turn allows us to calculate an arbitrary $\text{LCE}(i, j)$ efficiently. The key difference is that instead of calculating $\min(\text{LCE}(i, j), t)$ naively, we use a recursive approach. The main tool there is an efficient Union-Find structure. This is enough to answer $\mathcal{O}(n)$ short queries in $\mathcal{O}(n \log \log n \cdot \alpha(n \log \log n, n \log \log n))$ total time. We can remove the $\alpha(n \log \log n, n \log \log n)$ factor introducing another difference cover and carefully analyzing the running time of the Union-Find data structure. Finally, we modify the algorithm to work faster when the number of queries q is smaller than n . The main insight allowing us to obtain $\mathcal{O}(q \log \log n + n \log^* n)$ total time is introducing multiple levels of difference covers with some additional properties. Such family of difference covers was implicitly provided in [8].

2 Preliminaries

2.1 t -covers

A difference cover is a number-theoretic tool used throughout the paper. A set $\mathbf{D} \subseteq [0..t-1]$ is said to be a t -difference-cover if $[0..t-1] = \{(x - y) \bmod t : x, y \in \mathbf{D}\}$.

Lemma 2.1 (Maekawa [17]). *For every integer t there is t -difference-cover of size $\mathcal{O}(\sqrt{t})$, which can be constructed in $\mathcal{O}(\sqrt{t})$ time.*

A subset X of $[1..n]$ is t -periodic if for each $i \in [1..n-t]$ we have: $i \in X \Leftrightarrow i+t \in X$.

A set $\mathbf{S} \subseteq [1..n]$ is called a t -cover of $[1..n]$ if \mathbf{S} is t -periodic and there is a constant-time computable function h such that for $1 \leq i, j \leq n-t$ we have $0 \leq h(i, j) \leq t$ and $i+h(i, j), j+h(i, j) \in \mathbf{S}(t)$ (see Figure 1).

A t -cover can be obtained by taking a t -difference-cover \mathbf{D} and setting $\mathbf{S}(t) = \{i \in [1..n] : i \bmod t \in \mathbf{D}\}$. This is a well-known construction implicitly used in [7], for example.

Lemma 2.2. *For each $t \leq n$ there is a t -cover $\mathbf{S}(t)$ of size $\mathcal{O}(\frac{n}{\sqrt{t}})$ which can be constructed in $\mathcal{O}(\frac{n}{\sqrt{t}})$ time.*




Figure 1: An example of a 6-cover $\mathbf{S}(6) = \{2, 3, 5, 8, 9, 11, 14, 15, 19, 20, 21, 23\}$ (for $\mathbf{D} = \{2, 3, 5\}$), with the elements marked as black circles. For example, we have $h(3, 10) = 5$, since $3 + 5, 10 + 5 \in \mathbf{S}(6)$.

2.2 Disjoint-sets structure

Our another tool is a disjoint-sets data structure. In this problem we maintain a family of disjoint subsets of $[1..n]$, initially consisting of singleton sets. We perform Find queries asking for a subset containing a given element, and Union operations which merge two subsets.

Note that the extremely fast-growing Ackermann function [21] is defined for $i, j \in \mathbb{Z}_{>0}$ as

$$A(i, j) = \begin{cases} 2^j & \text{if } i = 1, \\ A(i-1, 2) & \text{if } i > 1 \text{ and } j = 1, \\ A(i-1, A(i, j-1)) & \text{if } i > 1 \text{ and } j > 1. \end{cases}$$

Moreover, for $n, m \in \mathbb{Z}_{>0}$ ($m \geq n$) one defines $\alpha(m, n) = \min\{i \geq 1 : A(i, \lfloor \frac{m}{n} \rfloor) > \log n\}$.

Lemma 2.3 (Tarjan [20]). *A sequence of up to n Union and m Find operations on an n -element set can be executed on-line in $\mathcal{O}(n + m \cdot \alpha(m + n, n))$ total time.*

Lemma 2.4. *For every $n, m \in \mathbb{Z}_{>0}$, we have $n + m \cdot \alpha(m + n, n) = \mathcal{O}(m + n \log^* n)$.*

Proof. First, observe that the Ackermann function $A(i, j)$ is monotone with respect to both coordinates and that $A(i, j) \geq 2^{i+j-1}$. These properties are easy to show by induction. Additionally, let us recall the row inverse of the Ackermann function is defined for $i, n \in \mathbb{Z}_{>0}$ as $a(i, n) = \min\{j \geq 1 : A(i, j) > \log n\}$. Note that $\alpha(m, n) \leq i$ if $\frac{m}{n} \geq a(i, n)$, in particular, $\alpha(m, n) \leq 2$ if $m \leq n \log^* n$; see [21]. We shall make two claims relating the $\alpha(m, n)$ and $a(i, n)$ functions. First, $\alpha(n, n) \leq 4 + a(3, n)$ for every $n \in \mathbb{Z}_{>0}$. This follows from:

$$\begin{aligned} A(4 + a(3, n), 1) &= A(3 + a(3, n), 2) = A(2 + a(3, n), A(3 + a(3, n), 1)) \geq \\ &A(3, 2^{3+a(3, n)}) \geq A(3, a(3, n)) > \log n. \end{aligned}$$

Moreover, the fact that $a(3, n) \leq 2 + \lceil \log a(2, n) \rceil$ for every $n \in \mathbb{Z}_{>0}$ is a consequence of:

$$\begin{aligned} A(3, 2 + \lceil \log a(2, n) \rceil) &= A(2, A(3, 1 + \lceil \log a(2, n) \rceil)) \geq A(2, 2^{3+\lceil \log a(2, n) \rceil}) \geq \\ &A(2, 8a(2, n)) \geq A(2, a(2, n)) > \log n. \end{aligned}$$

To prove the lemma, we consider two cases: If $m + n \geq n \cdot a(3, n)$, then $m \cdot \alpha(m + n, n) \leq 3m$, so $\mathcal{O}(n + m \cdot \alpha(m + n, n)) = \mathcal{O}(n + m)$. Otherwise, the claims that we made above imply:

$$m \cdot \alpha(m + n, n) \leq m \cdot \alpha(n, n) \leq n \cdot a(3, n) \cdot (4 + a(3, n)) \leq n(2 + \lceil \log a(2, n) \rceil)(6 + \lceil \log a(2, n) \rceil),$$

i.e., $n + m \cdot \alpha(m + n, n) = \mathcal{O}(n \log^2 a(2, n)) = \mathcal{O}(n \cdot a(2, n)) = \mathcal{O}(n \log^* n)$, as desired. \square

3 Generic LCE algorithm for general ordered alphabets

We define t -short LCE queries by restricting the answer to at most t :

$$\text{ShortLCE}_t(i, j) = \min(\text{LCE}(i, j), t).$$

We define a t -block as a fragment of the input text w which starts in $\mathbf{S}(t)$ and has length t . If a position in $\mathbf{S}(t)$ lies near the end of w , we form a t -block from a suffix of w and enough dummy symbols to reach length t . We also introduce t -coarse LCE queries, which are LCE queries restricted to positions from $\mathbf{S}(t)$ returning the number of matching t -blocks:

$$\text{CoarseLCE}_t(i, j) = \begin{cases} \lfloor \text{LCE}(i, j)/t \rfloor & \text{if } i, j \in \mathbf{S}(t), \\ \perp & \text{otherwise.} \end{cases}$$

We now describe how to use ShortLCE and CoarseLCE queries for general LCE queries.

Lemma 3.1. *If every sequence of q ShortLCE $_t$ queries and CoarseLCE $_t$ queries can be executed on-line in total time $T(n, q)$, then every sequence of q LCE queries can be executed on-line in total time $T(n, \mathcal{O}(q)) + \mathcal{O}(n + q)$.*

Proof. To calculate $\text{LCE}(i, j)$ we first check if $\text{LCE}(i, j) < t$ by calling $\text{ShortLCE}_t(i, j)$. If so, we are done. Otherwise, we can reduce computing $\text{LCE}(i, j)$ to computing $\text{LCE}(i + \Delta, j + \Delta)$ for any $\Delta \leq t$. In particular, we can choose $\Delta = h_t(i, j)$ so that $i + \Delta, j + \Delta \in \mathbf{S}(t)$. Then we call $\text{CoarseLCE}_t(i + \Delta, j + \Delta)$ which gives us the value $\lfloor \frac{1}{t}(\text{LCE}(i, j) - \Delta) \rfloor$. Computing the exact value of $\text{LCE}(i, j)$ requires another ShortLCE $_t$ query; see Algorithm 1. The whole process is illustrated in Figure 2. \square

Algorithm 1: GenericLCE(i, j)

```

 $\ell_1 = \text{ShortLCE}_t(i, j)$ 
if  $\ell_1 < t$  then return  $\ell_1$ 
 $\Delta = h_t(i, j)$   $\triangleright i + \Delta, j + \Delta \in \mathbf{S}(t)$ 
 $\ell_2 = t \cdot \text{CoarseLCE}_t(i + \Delta, j + \Delta)$ 
 $\ell_3 = \text{ShortLCE}_t(i + \Delta + \ell_2, j + \Delta + \ell_2)$ 
return  $\Delta + \ell_2 + \ell_3$ 

```




Figure 2: Illustration of Algorithm 1 for the case $\ell_1 \geq \Delta$.

4 ShortLCE $_t$ queries in $\mathcal{O}(\log t)$ amortized time

In this section we show how to implement fast on-line ShortLCE $_t$ queries. We assume that $t = 2^k$ and set $t' = \Theta(\log t)$ to be a smaller power of two. The amortized running time is $\mathcal{O}(\log t + \sqrt{\log t \log^* n})$, which in particular is $\mathcal{O}(\log t)$ for $t = \log^{\Omega(1)} n$. The key components are Union-Find structures and t' -covers. We start with a simpler (and slightly slower) algorithm without t' -covers.

4.1 ShortLCE_t queries in $\mathcal{O}(\log t \cdot \alpha((n+q) \log t, n \log t))$ amortized time

Lemma 4.1. *A sequence of q ShortLCE_{2^k}(i, j) queries can be executed on-line in total time $\mathcal{O}((q+n)k \cdot \alpha((q+n)k, nk))$.*

Proof. We compute ShortLCE_{2^k}(i, j) using a recursive procedure; see Algorithm 2. The procedure first checks if $w[i..i + 2^k - 1]$ is already known to be equal to $w[j..j + 2^k - 1]$ using a Union-Find structure. If so, we are done. Otherwise, if $k = 0$, we simply compare $w[i]$ and $w[j]$. If $k > 0$, we recursively calculate ShortLCE_{2^{k-1}}(i, j) and, if the call returns 2^{k-1} , also ShortLCE_{2^{k-1}}(i, j). Finally, if both calls return 2^{k-1} , we update the Union-Find structure to store that $w[i..i + 2^k - 1] = w[j..j + 2^k - 1]$.

Algorithm 2: ShortLCE_{2^k}(i, j): compute LCE(i, j) up to length 2^k

```

if Findk( $i$ ) = Findk( $j$ ) then return  $2^k$ 
if  $k = 0$  then
    if  $w[i] = w[j]$  then  $\ell = 1$  else  $\ell = 0$ 
else
     $\ell = \text{ShortLCE}_{2^{k-1}}(i, j)$ 
    if  $\ell = 2^{k-1}$  then
         $\ell = 2^{k-1} + \text{ShortLCE}_{2^{k-1}}(i + 2^{k-1}, j + 2^{k-1})$ 
if  $\ell = 2^k$  then Unionk( $i, j$ )
return  $\ell$ 

```

To analyze the complexity of the procedure, we first observe that the total number of calls to Union is $\mathcal{O}(nk)$, because each such call discovers that $w[i..i + 2^k - 1] = w[j..j + 2^k - 1]$ (which was not known before). Moreover, these calls contribute $\mathcal{O}(nk)$ to the total running time. We argue that the number of executed Find queries and the running time of the remaining operations performed by ShortLCE_{2^k}(i, j) is proportional to $\mathcal{O}(k+1)$ plus the number of Union calls, which implies the lemma. For the sake of conciseness, #union denotes the number of calls to Union triggered by the considered call to ShortLCE (including itself).

We inductively bound the number of recursive calls triggered by ShortLCE_{2^k}(i, j):

$$\begin{array}{ll}
 2k + 1 + 2\#\text{union} & \text{if } w[i..i + 2^k - 1] \neq w[j..j + 2^k - 1], \\
 1 + 2\#\text{union} & \text{if } w[i..i + 2^k - 1] = w[j..j + 2^k - 1].
 \end{array}$$

ShortLCE₁ terminates immediately, so this holds for $k = 0$. For $k > 0$ we have four cases.

1. $w[i..i + 2^k - 1]$ is already known to be equal to $w[j..j + 2^k - 1]$. Then we terminate immediately.
2. $w[i..i + 2^{k-1} - 1] \neq w[j..j + 2^{k-1} - 1]$. Then the number of recursive calls triggered by ShortLCE_{2^{k-1}}(i, j) is $2k - 1 + 2\#\text{union}$ so the number of recursive calls triggered by ShortLCE_{2^k}(i, j) is $2k + 2\#\text{union}$.
3. $w[i..i + 2^{k-1} - 1] = w[j..j + 2^{k-1} - 1]$ but $w[i + 2^{k-1}..i + 2^k - 1] \neq w[j + 2^{k-1}..j + 2^k - 1]$. The number of recursive calls triggered by ShortLCE_{2^{k-1}}(i, j) and ShortLCE_{2^{k-1}}($i + 2^{k-1}, j + 2^{k-1}$) is $1 + 2\#\text{union}$ and $2k - 1 + 2\#\text{union}$, respectively. The total number of triggered recursive calls is hence $2k + 1 + 2\#\text{union}$.
4. $w[i..i + 2^{k-1} - 1] = w[j..j + 2^{k-1} - 1]$ and $w[i + 2^{k-1}..i + 2^k - 1] = w[j + 2^{k-1}..j + 2^k - 1]$. The number of recursive calls triggered by both ShortLCE_{2^{k-1}}(i, j) and ShortLCE_{2^{k-1}}($i + 2^{k-1}, j + 2^{k-1}$) is $1 + 2\#\text{union}$. However, $w[i..i + 2^k - 1]$ was not known to be equal to $w[j..j + 2^k - 1]$, so we then execute Union_k(i, j). Hence the total number of recursive calls is $1 + 2\#\text{union}$ (rather than of $3 + 2\#\text{union}$).

Consequently, the total running time follows from Lemma 2.3. \square




Figure 3: A recursion tree of $\text{SparseShortLCE}_{t,t'}(i, j)$ for some example parameters such that $t = 2^4 t'$. The calls terminating with Union, Find and naive tests (in a segment of size t') are shown as nodes in the figure. The naive tests are only at the bottom of the tree and they are accompanied by Unions (except the last one).

4.2 Faster ShortLCE_t queries

Assume $t = 2^k = \Omega(\log n)$. We show how to reduce the factor $\alpha(qk + nk, nk)$ introducing a t' -cover, for $t' = 2^{k'}$. We define a sparse version of ShortLCE queries, which are ShortLCE queries restricted to positions from $\mathbf{S}(t')$:

$$\text{SparseShortLCE}_{t,t'}(i, j) = \begin{cases} \text{ShortLCE}_t(i, j) & \text{if } i, j \in \mathbf{S}(t') \\ \perp & \text{otherwise} \end{cases}$$

We slightly modify Algorithm 2 to obtain Algorithm 3, which computes $\min(\text{LCE}(i, j), 2^k)$ for positions $i, j \in \mathbf{S}(t')$.

Lemma 4.2. *A sequence of q $\text{SparseShortLCE}_{2^k, 2^{k'}}$ queries can be executed on-line in total time $\mathcal{O}(q(k + 2^{k'}) + n\sqrt{2^{k'}} + \frac{nk}{\sqrt{2^{k'}}} \log^* n)$.*

Proof. The analysis is similar to the proof of Lemma 4.1. The total number of calls to Union is now only $\mathcal{O}(\frac{nk}{2^{k'/2}})$ because we always have that $i, j \in \mathbf{S}(2^{k'})$. Hence, excluding the cost of computing $\ell = \text{ShortLCE}_{2^{k'}}(i, j)$, the total time complexity is $\mathcal{O}(qk + \frac{nk}{2^{k'/2}} \log^* n)$ by the same reasoning as in Lemma 4.1, except that we additionally apply Lemma 2.4 to bound the running time of the Union-Find data structure (stated in Lemma 2.3).

Algorithm 3: $\text{SparseShortLCE}_{2^k, 2^{k'}}(i, j)$: compute $\min(\text{LCE}(i, j), 2^k)$ for $i, j \in \mathbf{S}(2^{k'})$

```

if Findk(i) = Findk(j) then return 2k
if k = k' then
    Compute naively  $\ell = \text{ShortLCE}_{2^{k'}}(i, j)$ 
else
     $\ell = \text{SparseShortLCE}_{2^{k-1}, 2^{k'}}(i, j)$ 
    if  $\ell = 2^{k-1}$  then
         $\ell = 2^{k-1} + \text{SparseShortLCE}_{2^{k-1}, 2^{k'}}(i + 2^{k-1}, j + 2^{k-1})$ 
if  $\ell = 2^k$  then Unionk(i, j)
return  $\ell$ 

```

Algorithm 4: FasterShortLCE_{2^k, 2^{k'}}(*i*, *j*)

Compute naively $\ell = \text{ShortLCE}_{2^{k'}}(i, j)$
if $\ell < 2^{k'}$ **then return** ℓ
 $\Delta = h_{2^{k'}}(i, j)$
 $\ell = \Delta + \text{SparseShortLCE}_{2^k, 2^{k'}}(i + \Delta, j + \Delta)$
return $\min(\ell, 2^{k'})$

Now we analyze the cost of computing $\ell = \text{ShortLCE}_{2^{k'}}(i, j)$. First, observe that for every original call to $\text{SparseShortLCE}_{2^k, 2^{k'}}(i, j)$ we have at most one such computation with $\ell < 2^{k'}$ (because it means that we have found a mismatch and no further recursive calls are necessary). On the other hand, if $\ell = 2^{k'}$, then we call $\text{Union}_{k'}(i, j)$, which may happen at most $\frac{n}{2^{k'/2}}$ times. Therefore, the total complexity of all these naive computations is $\mathcal{O}(n2^{k'/2} + q \cdot 2^{k'})$. \square

The next lemma is a direct consequence of Lemma 4.2 and Algorithm 4 with $2^{k'} = \Theta(k)$.

Lemma 4.3. *A sequence of q ShortLCE_{2^k} queries can be executed on-line in total time $\mathcal{O}(qk + n\sqrt{k} \log^* n)$.*

5 CoarseLCE_t queries

Let $t = \Omega(\log^2 n)$. Recall that we defined a t -block of w as a factor of size t starting in $\mathbf{S}(t)$. We want to show how to preprocess w in $\mathcal{O}(n \log \log n)$ time, so that any CoarseLCE_t query can be answered in constant time. To this end we proceed as follows:

1. sort all t -blocks in lexicographic order and remove duplicates,
2. encode every t -block with its rank on the sorted list,
3. construct a new string $\text{code}(w)$ of length $\mathcal{O}(n)$ over alphabet $[1..n]$, such that any CoarseLCE_t query can be reduced to an LCE query on $\text{code}(w)$,
4. preprocess $\text{code}(w)$ for LCE queries.

Lemma 5.1. *For $t = \Omega(\log^2 n)$ we can lexicographically sort all t -blocks of w in $\mathcal{O}(n \log t)$ time.*

Proof. Two t -blocks can be lexicographically compared with a ShortLCE_t query. We have $\mathcal{O}(\frac{n}{\sqrt{t}})$ such blocks, hence one of the classical sorting algorithms they can be all sorted using $\mathcal{O}(\frac{n}{\sqrt{t}} \log n) = \mathcal{O}(n)$ queries. By Lemma 4.3, the total time to execute these queries and sort all t -blocks is therefore $\mathcal{O}(n \log t)$. \square

We can use the lexicographic order of t -blocks to assign ranks to all t -blocks. Then we reduce CoarseLCE queries to LCE queries in a word $\text{code}(w)$ over an integer alphabet; see Figure 4.

Lemma 5.2. *For $t = \Omega(\log^2 n)$ we can preprocess w in $\mathcal{O}(n \log t)$ time so that any CoarseLCE_t query can be answered in constant time.*

Proof. Using Lemma 5.1, we assign a number to each t -block, so that two t -blocks are identical if and only if their numbers are equal. The number assigned to the block starting at position $p \in \mathbf{S}(t)$ is denoted $\text{rank}(p)$. These numbers are ranks on a sorted list of length $|\mathbf{S}(t)|$, so $\text{rank}(p) \in [1..|\mathbf{S}(t)|]$. Then we construct a new string $\text{code}(w)$ as follows. Let $\{i_1, i_2, \dots, i_k\} = [1, t] \cap \mathbf{S}(t)$ and z_s be the word obtained from w by concatenating the numbers assigned to all t -blocks starting at positions $i_s, i_s + t, i_s + 2t, i_s + 3t, \dots$:

$$z_s = \text{rank}(i_s)\text{rank}(i_s + t)\text{rank}(i_s + 2t)\text{rank}(i_s + 3t)\dots$$




Figure 4: 6-blocks of w are lexicographically sorted (using ShortLCE_t) and ranked. Then $\text{CoarseLCE}_6(2, 11)$ in w is reduced to $\text{LCE}(1, 12)$ in $\text{code}(w)$.

Finally, we introduce k new distinct letters $\#_1, \#_2, \dots, \#_s$ and construct $\text{code}(w)$:

$$\text{code}(w) = z_1 \cdot \#_1 \cdot z_2 \cdot \#_2 \cdot z_3 \cdot \#_3 \cdots z_k \cdot \#_k.$$

Next, $\text{code}(w)$ is preprocessed to answer LCE queries in constant time. A $\text{CoarseLCE}_t(p, q)$ query for positions $p, q \in \mathbf{S}(t)$ is answered by first computing positions p', q' corresponding to p, q in $\text{code}(w)$. Formally, if $p = i_s \bmod t$, then $p' = |z_1 \#_1 z_2 \#_2 \dots z_{s-1} \#_{s-1}| + \frac{p - i_s}{t} + 1$; q' is computed similarly. Then an $\text{LCE}(p', q')$ query on $\text{code}(w)$ returns $\text{CoarseLCE}_t(p, q)$. The positions p' and q' can be computed in constant time, so the total query time is constant. Preprocessing $\text{code}(w)$ requires constructing its suffix array, which takes linear time for integer alphabets of polynomial size, and preprocessing it for range minimum queries, which also takes linear time. Hence the total preprocessing time is $\mathcal{O}(n \log t)$. \square

Theorem 5.3. *A sequence of $\mathcal{O}(n)$ LCE queries for a string over a general ordered alphabet can be executed on-line in total time $\mathcal{O}(n \log \log n)$ making only $\mathcal{O}(n)$ symbol comparisons.*

Proof. We set $t = \Theta(\log^2 n)$ and reduce each LCE query to constant number of CoarseLCE_t queries and ShortLCE_t queries as described in Lemma 3.1. Thus together with Lemma 4.3 and Lemma 5.2 we obtain that any sequence of q LCE queries for a string over a general ordered alphabet can be realized in $\mathcal{O}(n \log \log n)$ time. However, the total number of symbol comparisons used by the algorithm might be $\Omega(n \log \log n)$. This can be decreased to $\mathcal{O}(n)$ with yet another Union-Find data structure, where we maintain sets of positions already known to store the same letter. This is essentially the idea used in Lemma 7 of [12]. \square

6 Faster solution for sublinear number of queries

The algorithm presented in the previous section is not efficient when the number of queries q is significantly smaller than the length of the string n . In this section we show that this can be avoided, and we present an $\mathcal{O}(q \log \log n + n \log^* n)$ -time algorithm. This requires some nontrivial changes in our approach. In particular, we need a stronger notion of t -covers, which form a *monotone family*.

$\mathbf{S}(4^0), \mathbf{S}(4^1), \mathbf{S}(4^2), \dots \subseteq [1, n]$ is a monotone family of covers if the following conditions hold for every k :

1. $\mathbf{S}(4^k)$ is a 4^k -cover (except that h_{4^k} is computable in $\mathcal{O}(k)$ instead of constant time).
2. $\mathbf{S}(4^{k+1}) \subseteq \mathbf{S}(4^k)$.
3. For any $i, j \in \mathbf{S}(4^k)$ we have that $h_{4^{k+1}}(i, j) \in \{0, 4^k, 2 \cdot 4^k\}$, and furthermore for such arguments $h_{4^{k+1}}$ can be evaluated in constant time.
4. $|\mathbf{S}(4^k)| \leq (\frac{3}{4})^k n$.

The existence of such a family is not completely trivial, in particular plugging in the standard construction of $\mathbf{S}(4^k)$ from Lemma 2.1 does not guarantee that $\mathbf{S}(4^{k+1}) \subseteq \mathbf{S}(4^k)$. The following lemma, implicitly shown in [8], provides an efficient construction.

Lemma 6.1 (Gawrychowski et al. [8], Section 4.1). *Let $\mathbf{S}(4^k)$ be the set of non-negative integers $i \in [1, n]$ such that none of the k least significant digits of the base-4 representation of i is zero. Then $\mathbf{S}(4^0), \mathbf{S}(4^1), \mathbf{S}(4^2), \dots$ is a monotone family of covers, which can be constructed in $\mathcal{O}(n)$ total time.*

6.1 ShortLCE_t queries with monotone family of covers

Similarly as in the proof of Lemma 4.3, we reduce ShortLCE queries to SparseShortLCE queries. However, now we slightly change the definition of SparseShortLCE queries so that there is only one parameter as follows:

$$\text{SparseShortLCE}_t(i, j) = \begin{cases} \text{ShortLCE}_t(i, j) & \text{if } i, j \in \mathbf{S}(t) \\ \perp & \text{otherwise} \end{cases}$$

Lemma 6.2. *Consider a sequence of q SparseShortLCE_{4^{k_i}} queries for $i \in \{1, \dots, q\}$. The queries can be answered online in $\mathcal{O}((n+s) \cdot \alpha(n+s, n))$ time where $s = \sum_{i=1}^q T_i$ with $T_i = 1$ if the i -th query returns 4^{k_i} and $T_i = k_i + 1$ otherwise.*

Proof. We maintain a separate Union-Find structure for $\mathbf{S}(4^k)$ at every level $k \in \{0, \dots, K\}$ where $K = \max_{i=1}^q k_i$. To answer a query for SparseShortLCE_{4^k}, we check if $\text{Find}_k(i) = \text{Find}_k(j)$ and if so, return 4^k . Otherwise, we calculate the answer with at most four calls to SparseShortLCE_{4^{k-1}}. This is possible because $\mathbf{S}(4^k) \subseteq \mathbf{S}(4^{k-1})$ and $\mathbf{S}(4^{k-1})$ is 4^{k-1} -periodic. Finally, we call $\text{Union}_k(i, j)$ if the answer is 4^k ; see Algorithm 5.

Algorithm 5: SparseShortLCE_{4^k}(i, j): compute $\min(\text{LCE}(i, j), 4^k)$ for $i, j \in \mathbf{S}(4^k)$

```

if Findk( $i$ ) = Findk( $j$ ) then return  $4^k$ 
if  $k = 0$  then
    if  $w[i] = w[j]$  then  $\ell = 1$  else  $\ell = 0$ 
else
     $\ell = 0$ 
    for  $p = 0$  to 3 do
         $\ell = \ell + \text{SparseShortLCE}_{4^{k-1}}(i + p \cdot 4^{k-1}, j + p \cdot 4^{k-1})$ 
        if  $\ell < (p+1) \cdot 4^{k-1}$  then break
if  $\ell = 4^k$  then Unionk( $i, j$ )
return  $\ell$ 

```

We again analyze the number of recursive calls to SparseShortLCE_{4^k} counting Union operations. The total number of unions at level k is $|\mathbf{S}(4^k)| \leq (\frac{3}{4})^k$, and in total this sums up to $\mathcal{O}(n)$. The amortized number of Find queries executed by a call to SparseShortLCE_{4^k} is constant if $\text{LCE}(i, j) = 4^k$ and $\mathcal{O}(k+1)$ otherwise. These values also bound the running time of the remaining operations. Hence, by Lemma 2.3, the total time is as claimed. \square

Lemma 6.3. *A sequence of q queries ShortLCE_{4^{k_i}} for $i \in \{1, \dots, q\}$ can be answered online in total time $\mathcal{O}((n+s) \cdot \alpha(n+s, n)) = \mathcal{O}(n \log^* n + s)$ where $s = \sum_{i=1}^q (k_i + 1)$.*

Proof. We calculate ShortLCE_{4^k}(i, j) using $\mathcal{O}(k)$ SparseShortLCE queries; see Algorithm 6. We iterate through $k' = 0, 1, \dots, k-1$ maintaining Δ such that $0 \leq \Delta \leq \text{LCE}(i, j)$ and $i + \Delta, j + \Delta \in \mathbf{S}(4^{k'})$. Before incrementing k' , we keep increasing Δ by $4^{k'}$ until $i + \Delta, j + \Delta \in \mathbf{S}(4^{k'})$ or $\Delta > \text{LCE}(i, j)$. The latter

k'	SparseShortLCE calls
0	$\text{SparseShortLCE}_{4^0}(10130_4, 00101_4) \rightarrow \Delta = 00001_4$
1	$\text{SparseShortLCE}_{4^1}(10131_4, 00102_4) \rightarrow \Delta = 00011_4$
1	$\text{SparseShortLCE}_{4^1}(10201_4, 00112_4) \rightarrow \Delta = 00021_4$
3	$\text{SparseShortLCE}_{4^3}(10211_4, 00122_4) \rightarrow \Delta = 01021_4$
return call	$\text{SparseShortLCE}_{4^4}(11211_4, 01122_4)$

Figure 5: An execution of $\text{ShortLCE}_{4^4}(i = (10130)_4, j = (00101)_4)$ (assuming $\text{LCE}(i, j) > 4^4$). The numbers are given in base-4 representation. Note that there is no $\text{SparseShortLCE}_{4^2}$ call.

condition is checked by calling $\text{SparseShortLCE}_{4^{k'}}(i + \Delta, j + \Delta)$ and terminating if it returns less than $4^{k'}$. The while loop iterates at most twice, because $h_{4^{k'+1}} \in \{0, 4^{k'}, 2 \cdot 4^{k'}\}$. Eventually, we either terminate having found the answer, or we can obtain it with a single call to $\text{SparseShortLCE}_{4^k}(i + \Delta, j + \Delta)$.

Algorithm 6: $\text{ShortLCE}_{4^k}(i, j)$

```

 $\ell = \Delta = 0$ 
for  $k' = 0$  to  $k - 1$  do
  while  $i + \Delta \notin \mathbf{S}(4^{k'+1})$  or  $j + \Delta \notin \mathbf{S}(4^{k'+1})$  do
     $\ell = \ell + \text{SparseShortLCE}_{4^{k'}}(i + \Delta, j + \Delta)$   $\triangleright i + \Delta, j + \Delta \in \mathbf{S}(4^{k'})$ 
     $\Delta = \Delta + 4^{k'}$ 
  if  $\ell < \Delta$  then return  $\min(4^k, \ell)$ 
return  $\min(4^k, \Delta + \text{SparseShortLCE}_{4^k}(i + \Delta, j + \Delta))$   $\triangleright i + \Delta, j + \Delta \in \mathbf{S}(4^k)$ 

```

Let us analyze the total time complexity. Each call to ShortLCE_{4^k} performs up to k $\text{SparseShortLCE}_{4^{k'}}$ queries, but we terminate as soon as we obtain an answer other than $4^{k'}$. In Lemma 6.2, the last of these queries contributes $\mathcal{O}(k' + 1) = \mathcal{O}(k + 1)$ to s , while the remaining queries contribute one each. The total contribution of all $\text{SparseShortLCE}_{4^{k'}}$ queries called by a single ShortLCE_{4^k} query is therefore $\mathcal{O}(k + 1)$. Hence, the total running time consumed by all $\text{SparseShortLCE}_{4^{k'}}$ queries is $\mathcal{O}((n + s) \cdot \alpha(n + s, n))$ where $s = \mathcal{O}(\sum_{i=1}^q (k_i + 1))$. It is not hard to see that the remaining time consumed by a single ShortLCE_{4^k} query is $\mathcal{O}(k + 1)$. This is partly because checking whether $i + \Delta$ and $j + \Delta$ belong to $\mathbf{S}(4^{k'+1})$ takes constant time, since we know that these indices are in $\mathbf{S}(4^{k'})$. Over all queries this sums up to $\mathcal{O}(s)$, which is dominated by the running time of the $\text{SparseShortLCE}_{4^{k'}}$ queries. The $\mathcal{O}(n \log^* n + s)$ upper bound follows from Lemma 2.4. \square

6.2 Final algorithm

We first modify the implementation details for CoarseLCE to reduce the preprocessing time.

Lemma 6.4. *For $t = \Omega(\log^6 n)$ we can preprocess a string of length n in $\mathcal{O}(n \log^* n)$ time, so that each CoarseLCE_t query can be answered in constant time.*

Proof. We set $k = \lceil \frac{1}{2} \log t \rceil$ and lexicographically sort all 4^k -blocks using ShortLCE_{4^k} queries. The number of blocks is at most $(\frac{3}{4})^k n \leq \frac{n}{t^{0.5 \log 0.75}} \leq \frac{n}{t^{0.2}}$. By Lemma 6.3, the sorting time is:

$$\mathcal{O}\left(\frac{n}{t^{0.2}} \log n \log t + n \log^* n\right) = \mathcal{O}\left(n \frac{\log n \log \log n}{\log^{1.2} n} + n \log^* n\right) = \mathcal{O}(n \log^* n).$$

Then we proceed as in the proof of Lemma 5.2. \square

By combining Lemmas 6.3 and 6.4, we obtain the final theorem.

Theorem 6.5. *A sequence of q LCE queries for a string over a general ordered alphabet can be executed on-line in total time $\mathcal{O}(q \log \log n + n \log^* n)$ making $\mathcal{O}(q + n)$ symbol comparisons.*

7 Final remarks

We gave an $\mathcal{O}(n \log \log n)$ -time algorithm for answering on-line $\mathcal{O}(n)$ LCE queries for general ordered alphabet. It is known (see [14]) that the runs of the string can be computed in $\mathcal{O}(T(n))$ time, where $T(n)$ is the time to execute on-line $\mathcal{O}(n)$ LCE queries. Hence our algorithm implies the following result:

Corollary 7.1. *The runs of a string over general ordered alphabet can be computed in $\mathcal{O}(n \log \log n)$ time.*

Our algorithm is a major step towards a positive answer for a question posed by Kosolobov [14], who asked if $\mathcal{O}(n)$ time algorithm is possible.

It is also natural to consider general unordered alphabets, that is, strings where the only allowed operation is checking equality of two characters.

Theorem 7.2. *A sequence of q LCE queries for a string over a general unordered alphabet can be executed in $\mathcal{O}(q \log n + n \log^* n)$ time making $\mathcal{O}(n + q)$ symbol equality-tests.*

Proof. We can use the faster ShortLCE_{4^k} algorithm described in Section 6.1 with $k = \lceil \frac{1}{2} \log n \rceil$. Observe that in this approach we did not use the order of the characters, and thus it still works for unordered alphabets. \square

Note that for unordered alphabets the reduction by Kosolobov [14] (see also [2]) from computing runs to LCE queries no longer works. Actually, deciding whether a given string is square-free already requires $\Omega(n \log n)$ comparisons, as shown by Main and Lorentz [18]. On the other hand for $\mathcal{O}(n)$ LCE queries $\mathcal{O}(n)$ equality tests always suffice.

References

- [1] Hideo Bannai, Tomohiro I, Shunsuke Inenaga, Yuto Nakashima, Masayuki Takeda, and Kazuya Tsuruta. A new characterization of maximal repetitions by Lyndon trees. In Piotr Indyk, editor, *26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 562–571. SIAM, 2015. doi:10.1137/1.9781611973730.38.
- [2] Hideo Bannai, Tomohiro I, Shunsuke Inenaga, Yuto Nakashima, Masayuki Takeda, and Kazuya Tsuruta. The “runs” theorem, 2015. arXiv:1406.0263v7.
- [3] Philip Bille, Paweł Gawrychowski, Inge Li Gørtz, Gad M. Landau, and Oren Weimann. Longest common extensions in trees. *Theoretical Computer Science*, 2015. In press. doi:10.1016/j.tcs.2015.08.009.
- [4] Philip Bille, Inge Li Gørtz, Mathias Bæk Tejs Knudsen, Moshe Lewenstein, and Hjalte Wedel Vildhøj. Longest common extensions in sublinear space. In Ferdinando Cicalese, Ely Porat, and Ugo Vaccaro, editors, *Combinatorial Pattern Matching, CPM 2015*, volume 9133 of *LNCS*, pages 65–76. Springer, 2015. doi:10.1007/978-3-319-19929-0_6.
- [5] Philip Bille, Inge Li Gørtz, Benjamin Sach, and Hjalte Wedel Vildhøj. Time-space trade-offs for longest common extensions. *Journal of Discrete Algorithms*, 25:42–50, 2014. doi:10.1016/j.jda.2013.06.003.
- [6] Dany Breslauer. *Efficient String Algorithmics*. PhD thesis, Columbia University, 1992. URL: <http://www.cs.columbia.edu/~library/theses/breslauer.ps.gz>.

- [7] Stefan Burkhardt and Juha Kärkkäinen. Fast lightweight suffix array construction and checking. In Ricardo A. Baeza-Yates, Edgar Chávez, and Maxime Crochemore, editors, *Combinatorial Pattern Matching, CPM 2003*, volume 2676 of *LNCS*, pages 55–69. Springer, 2003. doi:10.1007/3-540-44888-8_5.
- [8] Paweł Gawrychowski, Tomasz Kociumaka, Jakub Radoszewski, Wojciech Rytter, and Tomasz Waleń. Universal reconstruction of a string. In Frank Dehne, Jörg-Rüdiger Sack, and Ulrike Stege, editors, *Algorithms and Data Structures, WADS 2015*, volume 9214 of *LNCS*, pages 386–397. Springer, 2015. doi:10.1007/978-3-319-21840-3_32.
- [9] Shunsuke Inenaga. A faster longest common extension algorithm on compressed strings and its applications. In Jan Holub and Jan Ždárek, editors, *Prague Stringology Conference 2015*, pages 1–4. Department of Theoretical Computer Science, Faculty of Information Technology, Czech Technical University in Prague, 2015. URL: <http://www.stringology.org/event/2015/p01.html>.
- [10] Marek Karpiński, Wojciech Rytter, and Ayumi Shinohara. An efficient pattern-matching algorithm for strings with short descriptions. *Nordic Journal of Computing*, 4(2):172–186, 1997.
- [11] Dmitry Kosolobov. Finding the leftmost critical factorization on unordered alphabet, 2015. arXiv:1509.01018.
- [12] Dmitry Kosolobov. Lempel-Ziv factorization may be harder than computing all runs. In Ernst W. Mayr and Nicolas Ollinger, editors, *Theoretical Aspects of Computer Science, STACS 2015*, volume 30 of *LIPIcs*, pages 582–593. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2015. doi:10.4230/LIPIcs.STACS.2015.582.
- [13] Dmitry Kosolobov. Online detection of repetitions with backtracking. In Ferdinando Cicalese, Ely Porat, and Ugo Vaccaro, editors, *Combinatorial Pattern Matching, CPM 2015*, volume 9133 of *LNCS*, pages 295–306. Springer, 2015. doi:10.1007/978-3-319-19929-0_25.
- [14] Dmitry Kosolobov. Computing runs on a general alphabet. *Information Processing Letters*, 116(3):241–244, 2016. doi:10.1016/j.ipl.2015.11.016.
- [15] Gad M. Landau and Uzi Vishkin. Fast parallel and serial approximate string matching. *Journal of Algorithms*, 10(2):157–169, 1989. doi:10.1016/0196-6774(89)90010-2.
- [16] Yury Lifshits. Processing compressed texts: A tractability border. In Bin Ma and Kaizhong Zhang, editors, *Combinatorial Pattern Matching, CPM 2007*, volume 4580 of *LNCS*, pages 228–240. Springer, 2007. doi:10.1007/978-3-540-73437-6_24.
- [17] Mamoru Maekawa. A \sqrt{n} algorithm for mutual exclusion in decentralized systems. *ACM Transactions on Computer Systems*, 3(2):145–159, May 1985. doi:10.1145/214438.214445.
- [18] Michael G. Main and Richard J. Lorentz. An $O(n \log n)$ algorithm for finding all repetitions in a string. *Journal of Algorithms*, 5(3):422–432, 1984. doi:10.1016/0196-6774(84)90021-X.
- [19] Masamichi Miyazaki, Ayumi Shinohara, and Masayuki Takeda. An improved pattern matching algorithm for strings in terms of straight-line programs. In Alberto Apostolico and Jotun Hein, editors, *Combinatorial Pattern Matching, CPM 1997*, volume 1264 of *LNCS*, pages 1–11. Springer, 1997. doi:10.1007/3-540-63220-4_45.
- [20] Robert E. Tarjan. Efficiency of a good but not linear set union algorithm. *Journal of the ACM*, 22(2):215–225, April 1975. doi:10.1145/321879.321884.
- [21] Robert E. Tarjan and Jan van Leeuwen. Worst-case analysis of set union algorithms. *Journal of the ACM*, 31(2):245–281, 1984. doi:10.1145/62.2160.